

### **REMARKS**

This amendment is responsive to the Office Action dated July 22, 2004. Applicant has amended claims 1, 11, 12, 23, 34, and 39 and cancelled claims 9, 20, 32, and 40. Claims 1-8, 10-19, 21-31, 33-39, and 41 are pending upon entry of this amendment.

### **Claim Objections**

In the Office Action, the Examiner objected to claim 11. Applicant has amended claim 11 for purposes of clarification.

### **Claim Rejection Under 35 U.S.C. § 102**

In the Office Action, the Examiner rejected claims 1-7, 9-18, 20-28, 32, 33, and 39-41 under 35 U.S.C. 102(e) as being anticipated by Nolan et al., U.S. Patent 6,640,278 (hereafter Nolan et al.). Applicant respectfully traverses the rejection to the extent such rejection may be considered applicable to the amended claims. Nolan et al. fails to disclose each and every feature of the claimed invention, as required by 35 U.S.C. 102(e), and provides no teaching that would have suggested the desirability modify the teachings of Nolan et al. to include such features.

#### *Independent claims 1, 12, and 23, 39 and 40*

For example, Nolan et al. fails to teach or suggest receiving a multicast message from a remote storage device, detecting the attachment of the remote storage device to a network in response to receiving the multicast message, and automatically incorporating storage capacity of the remote storage device into a logical storage volume, as recited by Applicant's independent claims 1, 12, 23, 39 and 40.

As described in the present application, a NAS master detects the connection of one or more new NAS storage devices to a network and automatically incorporates the additional storage capacity into one or more logical volumes. The NAS master incorporates the additional storage capacity according to various policies defined by a system administrator. For example, the system administrator may define one or more policies for dividing new storage capacity and allocating portions to different logical volumes. After incorporating the new storage device, the NAS master expands the file system based on the newly available storage area and adjusted

logical volumes. Automatically incorporating the additional storage capacity and expanding the file system accordingly advantageously allows a system administrator to simply connect a NAS storage device to a network when storage space is running low.

As described in detail in the present application, detecting attachment of a remote storage device to a network and automatically incorporating storage capacity of the remote storage device into a logical storage volume achieves certain technical advantages. For example, detecting attachment of a storage device to a network may facilitate the construction of a storage system that appears as seamless, single point of storage in which a single file can span multiple storage devices. As a result, a system administrator may be able to simply connect a storage device to a network when storage space is running low with the advantage that the additional storage capacity can be automatically incorporated without further intervention by the administrator.

In rejecting claims 1, 12, 23 and 39, the Examiner relies on Nolan et al. and refers to column 18 lines 5-22, column 2 lines 19-41, and column 34 lines 1-13. Specifically, the Examiner stated that Nolan et al. teaches a method wherein “storage devices are detected by a storage server and storage capacity is updated with the additional storage amount from the attached storage unit.” Contrary to the Examiner’s assertion Nolan et al. makes no mention of detecting the attachment of a remote storage device as required by independent claims 1, 12, 23, and 40. Rather, Nolan et al. instead refers to merely managing storage resources in a storage network according to storage domains. Specifically, Nolan et al. describes a “storage roll call table (that) includes a list of all the active storage devices detected by the storage server” (column 18 lines 12-14).

Applicant submits that detecting active storage devices, as described by Nolan, is fundamentally different than detecting the attachment of a new remote storage device, as recited in Applicant’s claims. An active storage device is inherently already connected to the network, while detecting the attachment of a storage device implies that the storage device was previously not connected to the network and therefore the storage capacity of the storage device not yet available to the network.

Moreover, Nolan fails to teach or suggest detecting the attachment of the remote storage device in response to receiving a multicast message, as required by amended claim 1. As

described in the present application, the NAS master detects a new NAS storage device by listening for multicast messages. In other words, the claimed invention may detect the attachment of a remote storage device by receiving a multicast message from the remote storage device.

Specifically, the Examiner relies on Nolan et al. and refers to column 29 line 61 – column 30 line 64 and states that Nolan et al. describes a method in which a “storage domain manager acts as a logical volume manager that responds to the detection and management of remote storage devices.” However, Applicant submits that Nolan et al. does not teach or even suggest detecting the attachment of the remote storage device comprises receiving a multicast message from the remote storage device. Rather, Nolan et al. instead refers to a hot copy request that contains the source member identifier, the destination member identifier, the copy block size, and the copy frequency or priority to perform a hot copy process that allows a data set to be copied from one individual member that is not yet down to a replacement drive. In order to clarify these differences, Applicant has amended independent claims 1, 11, and 23 to include the subject matter formerly recited in claims 9, 20, and 32 respectively.

Nolan et al. also fails to mention automatically incorporating storage capacity of the remote storage device into a logical storage volume as recited by claims 1, 12, 23, and 39. In fact, Nolan et al. fails to mention automatically incorporating storage capacity at all, much less automatically incorporating storage capacity upon detecting attachment of a storage device in a network. In general, Nolan et al. describes a method that allows configuration and management of current heterogeneous storage resources that simplifies management of storage systems. In particular, Nolan et al. describes a storage domain manager that receives storage transactions and manages storage transactions according to storage domains. The storage domain manager allows for flexible configuration, redundancy, failover, data migration, caching, and support of multiple protocols. In other words, Nolan et al. describes configuration and management of storage transactions, such as hot copy processes that migrate data between storage devices within the network.

Consequently, Nolan et al. does not teach or even suggest automatically incorporating storage capacity of a newly attached storage device into a logical storage volume, as required by

claims 1, 12, and 23, 39, and 40 and refers instead to managing storage transactions within storage devices previously attached to the network.

*Claims 2, 13, 24 and 41*

In rejecting claims 2, 13, 24, and 41 the Examiner relies on Nolan et al. and refers to column 10, lines 36-50. Specifically, the Examiner stated that Nolan et al. describes a method that “incorporates the storage into an existing storage volume.” Contrary to the Examiner’s assertion, Nolan et al. makes no mention of incorporating the storage capacity into an existing logical volume as recited by claims 2, 13, and 24 which are dependent on independent claims 1, 11, and 23, respectively. Rather, Nolan et al. refers simply to adding a new storage device to the network and migrating a data set to the new storage device. In particular, Nolan et al. describes a hot copy process that migrates data from one storage device to another which is fundamentally different from incorporating storage capacity into an existing logical storage volume. Logical storage volumes, as defined in the present application, “are groups of information located on physical volumes. Data on logical volumes appears to be contiguous to the user but can be discontinuous on the physical volume.” Consequently, it is clear that Nolan et al. does not teach or even suggest incorporating storage capacity into an existing logical volume, much less automatically incorporating additional storage capacity of a storage device upon detecting attachment of the storage device to the network.

*Claims 3, 14 and 25*

In rejecting claims 3, 14, and 25 the Examiner relies on Nolan et al. and refers to column 10, lines 39-50. Specifically the Examiner stated that Nolan et al. describes a method that “incorporates the storage into a new a new storage volume.” Contrary to the Examiner’s assertion, Nolan et al. makes no mention of incorporating the storage capacity into a new logical volume as recited by claims 3, 14, and 25 which are dependent on independent claims 1, 12, and 23, respectively. Rather, Nolan et al. instead refers to simply migrating a data set from an existing storage device to a new storage device and fails to mention incorporating storage capacity into a new logical volume. As argued previously, incorporating storage capacity into a new logical volume is fundamentally different from adding a new storage. Consequently, it is clear that Nolan et al. does not teach or even suggest incorporating storage capacity into a new

Application Number 09/808,313

Responsive to Office Action mailed July 22, 2004

logical volume, much less automatically incorporating additional storage capacity of a storage device upon detecting attachment of the storage device to the network.

*Claims 4, 15 and 26*

In rejecting claims 4, 15, and 26, the Examiner relies on Nolan et al. and refers to column 25, lines 14-25, and column 35, lines 3-25. Specifically, the Examiner stated that Nolan et al. describes a method in which “storage space automatically grows in response to a detected additional storage device.” Contrary to the Examiner’s assertion, Nolan et al. fails to mention automatically expanding a size of a storage space provided by a file system to include the storage capacity as recited by claims 4, 15, and 26 which are dependent on independent claims 1, 13, and 23, respectively. Rather, Nolan et al. instead simply refers to growing the storage capacity of a storage area network and a Storage Rollcall Table which lists all connected storage on a server. In contrast, the present application describes a NAS management module which interacts with a file system to expand the file system based on the newly available storage area and adjusted local volumes. For example, NAS management module may make one or more kernel call to the file system to increase the size of an existing storage space. In this manner, the NAS master automatically detects and incorporates storage capacity of a newly attached NAS storage device into one or more existing single points of storage in a file system. Consequently, Nolan et al. fails to teach or suggest automatically expanding a size of a storage space provided by a file system to include newly attached storage space.

*Claims 5, 16 and 27*

In rejecting claims 5, 16, and 27 the Examiner relies on Nolan et al. and refers to column 8, lines 24-35. Specifically, the Examiner stated that Nolan et al. describes a method in which “rule based management is used for managing storage capacity incorporation.” Contrary to the Examiner’s assertion, Nolan et al. fails to mention incorporating storage capacity according to one or more pre-defined policies as recited by claims 5, 16, and 27 which are dependent on independent claims 1, 13, and 23, respectively. Rather, Nolan et al. instead refers rules-based management of the system including scheduling, monitoring the system, informed consent management, and handling system processes and events. In general, Nolan et al. describes rule-based management for configuring and maintaining the storage server. However, the claimed invention is directed towards automatically incorporating storage capacity according to one or more pre-defined policies that specify the manner in which the storage capacity is to be incorporated into the file system. For example, one policy may specify that all storage is to be

added to a particular existing logical volume within the file system. Consequently, Nolan et al. fails to mention or even suggest incorporating storage capacity according to pre-defined policies as recited by claims 5, 16, and 27 much less automatically incorporating storage capacity of newly attached storage devices according to pre-defined policies.

*Claims 6, 17 and 28*

In rejecting claims 6, 17, and 28 the Examiner relies on Applicant's Admitted Prior Art (hereafter AAPA) and refers to page 1 lines 19-31 and column 2 lines 9-18 and 42-67 of the present application. Specifically the Examiner stated that Nolan et al. describes "storage devices [that] are integrated into the network where the storage domain is configured and comprises the storage area of the storage devices" and the AAPA discloses "assignment of a physical volume to a volume group."

Contrary to the Examiner's assertion, Nolan et al. fails to mention forming a network-based connection with the remote storage device, creating a physical volume for the remote storage device, and adding the physical volume to a volume group as recited by claims 6, 17, and 28 which are dependent on independent claims 1, 13, and 23, respectively. Rather, Nolan et al. instead refers to managing storage resources in a storage network according to a storage domain by assigning a logical storage extent to clients, mapping storage resources to the logical storage extents of the clients, and managing a storage domain of storage resources via a storage domain manager. In particular, Nolan et al. refers to a hot copy process that migrates data between storage devices in the network in a manner transparent to the clients rather than creating a physical volume for the remote storage device and adding the physical volume to a volume group. Consequently, Nolan et al. fails to teach or even suggest the necessary steps required to incorporate storage capacity as recited by claims 6, 17, and 28. Moreover, Applicant has not admitted that assignment of a physical volume to a volume group in response to detecting a new storage device by intercepting a multicast message is prior art.

*Claims 7, 18 and 29*

In rejecting claims 7, 18, and 29 the Examiner relies on AAPA page 1 lines 19-31 and column 8, lines 24-35. Specifically, the Examiner stated that Nolan et al. describes "rule-based management is used for managing storage capacity incorporation which could comprise selection of volume group" and the AAPA discloses "assignment of physical volume to a volume group."

Contrary to the Examiner's assertion, Nolan et al. fails to mention selecting the volume group from a number of volume groups based on one of the pre-defined policies as recited by claims 7, 18 and 29 which are dependent on independent claims 1, 13, and 23, respectively. As argued previously, Nolan et al. refers instead to rules-based management for configuring and maintaining the server, not managing storage capacity incorporation. Consequently, Nolan et al. fails to teach or even suggest selecting the volume group from a number of volume groups based on one of the pre-defined policies as required by claims 7, 18, and 29. Moreover, Applicant has not admitted that assignment of a physical volume to a volume group in response to detecting a new storage device by intercepting a multicast message is prior art.

*Claims 10 and 21*

In rejecting claims 10 and 21 the Examiner relies on Nolan et al. and refers to column 2, lines 9-67, and column 4, lines 63, – column 5, line 6. Specifically, the Examiner stated that Nolan et al. describes a “storage domain manager [that] acts as a logical volume manager that responds to the detection and management of remote storage devices.” Contrary to the Examiner's assertion, Nolan et al. fails to mention automatically controlling a logical volume manager in response to the detected remote storage device as recited by claims 10 and 21 which are dependent on independent claims 1 and 13. Rather, Nolan et al. instead refers to a storage domain manager that receives storage transactions and manages the received storage transactions according to configuration logic in the storage domain manager. As described in the present application, the NAS management module directs the logical volume manager to create a physical volume after detecting a new storage device and directs the logical volume manager to add the physical volume to an existing volume group based on pre-defined policies. As previously argued, Nolan et al. does not teach or even suggest automatically incorporating storage capacity of a newly attached storage device into one or more existing local storage spaces much less automatically controlling a logical volume manager in response to the detected remote storage device as recited in claims 10 and 21.

*Claims 11, 22 and 33*

In rejecting claims 11, 22, and 33 the Examiner relies on Nolan et al. and refers to column 5 line 6 – column 6 line 40, column 7 lines 41-61, and column 9 line 57 – column 10 line 7. Specifically, the Examiner stated that Nolan et al. describes a “connection of the storage device is



implemented via a network interface and the storage director manages the physical storage and adds to the physical resources of the storage domain. Contrary to the Examiner's assertion, Nolan et al. fails to mention instantiating a network driver to form a network-based connection with the remote storage device, directing the logical volume manager to create a physical volume for the remote storage device as if the remote storage device were local to the logical volume manager, and directing the logical volume manager to add the physical volume to a volume group as required by currently amended claim 11 and claims 22 and 33 which are dependent on independent claims 1, 13, and 23, respectively. Rather, Nolan et al. instead refers to a storage server including storage director logic that provides accurate configuration information and control by allowing automatic maintenance of the mapping of data sets in physical storage to servers. Automatic maintenance of the mapping of data sets in physical storage to servers is fundamentally different than automatically controlling the logical volume manager as specified in claims 11, 22, and 33. While mapping data sets may require instantiating a network drive to form a network-based connection and the storage director logic can configure servers as if the servers are local to the storage server, Nolan et al. fails to mention creating a physical volume for the remote storage device and directing the logical volume manager to add the physical volume to a volume group. Consequently, Nolan et al. does not teach or even suggest automatically controlling the logical volume manager in response to detecting a remote storage device as required by claims 11, 22, and 33.

Nolan et al. fails to disclose each and every limitation set forth in claims 1-7, 9-18, 20-28, 32, 33, 39, and 41 as amended. For at least these reasons, the Examiner has failed to establish a prima facie case for anticipation of Applicant's claims 1-7, 9-18, 20-28, 32, 33, 39, and 41 under 35 U.S.C. 102(e). Withdrawal of this rejection is requested.

#### **Claim Rejection Under 35 U.S.C. § 103**

In the Office Action, the Examiner rejected claims 8, 19, and 30 under 35 U.S.C. 103(a) as being unpatentable over Nolan et al. (US 6,640,278) in view of Carter et al. (US 5,987,506).

In the Office Action, the Examiner rejected claims 31-38 under 35 U.S.C. 103(a) as being unpatentable over Nolan et al. (US 6,640,278) in view of Lee et al. (US 6,601,101).

Applicant respectfully traverses the rejection. The applied references, alone or in

combination, fail to disclose or suggest the inventions defined by Applicant's claims, and provide no teaching that would have suggested the desirability of modification to arrive at the claimed invention.

*Claims 8, 19, and 30*

With respect to claims 8, 19, and 30, the Examiner proposes to modify the method of configuration and management of storage resources in a storage network described by Nolan et al. with the globally addressable storage environment described by Carter et al. Specifically, the Examiner refers to column 33 line 59 – column 34, line 30 and stated Carter et al. teaches “the exchange of IP address information upon the connection detection of storage device.” Carter et al., however, lacks any teaching or suggestion of any of the elements recited by claims 8, 19, and 30. Contrary to the Examiner's assertion, Carter et al. describes establishing a connection between networked computers that includes exchanging IP addresses rather than intercepting a request from the remote storage device for an IP address to detect attachment of the remote storage device to the network as required by claims 8, 19, and 30.

It is well established that the Examiner bears the burden of establishing a prima facie case of obviousness.<sup>1</sup> In doing so, the Examiner must determine whether the prior art provides a “teaching or suggestion to one of ordinary skill in the art to make the changes that would produce” the claimed invention.<sup>2</sup> A prima facie case of obviousness is established only when this burden is met.

As stated above, the Examiner's assertion that it would have been obvious to one of ordinary skill in the art to modify the configuration and management of storage resources of Nolan et al. with a globally addressable storage environment of Carter et al. is unreasonable and not supported by substantial evidence. As argued previously, Nolan et al. fails to mention or even suggest detecting the attachment of a remote storage device. Consequently, modification of the method for configuration and management of storage resources in a storage network of Nolan et al. with the globally addressable storage environment of Carter et al. would fail to achieve the Applicant's claimed invention. The Examiner's assertion that it would have been *prima facie* obvious to one skilled in the art to acquire an IP address for the remote device for the purpose of

---

<sup>1</sup> *In re Oetiker*, 24 USPQ2d 1443, 1445 (Fed. Cir. 1992).

<sup>2</sup> *In re Chu*, 36 USPQ2d 1089, 1094 (Fed. Cir. 1995).

device identification with respect to claims 8, 19, and 30 is unsupported. This assertion is flawed as neither Nolan et al. nor Carter et al. mention detecting the attachment of a remote storage device and the motivation is not found within the art, as required.

The conclusion of obviousness advanced by the Examiner relies on a motivation plucked directly from Applicant's own disclosure, rather than the prior art. Indeed, the Examiner cited no prior art teaching as the source for the motivation. Moreover, the Examiner failed to explain why one of ordinary skill in the art would have considered it desirable "to present the images for remote users as close as possible to the original and minimize customer complaints" in the Hess system.

The Court of Appeals for the Federal Circuit has established that motivation to combine references must be found in the prior art, and that it is impermissible hindsight for the Examiner to use the motivation stated in Applicant's own disclosure as a blueprint to reconstruct the claimed invention from the prior art.<sup>3</sup> Moreover, it is insufficient to merely pull such motivation out of thin air. Rather, the Examiner's rejection must be based on substantial evidence in the record demonstrated that the motivation for making the claimed invention resides in the prior art.<sup>4</sup> In summary, the Examiner's conclusion of obviousness is unsupported by evidence in the record.

With respect to claims 31-38, the Examiner proposes to modify the method for configuration and management of storage resources in a storage network described by Nolan et al. with the method for handing off TCP sessions in a system described by Lee et al. As argued previously, Nolan et al. fails to mention or even suggest detecting the attachment of a remote storage device. Consequently, modification of the method described by Nolan et al. with the method described by Lee et al. would fail to achieve the Applicant's claimed invention. Moreover, the Examiner has not pointed out any motivation in the art for such a modification.

#### *Claims 31 and 38*

In rejecting claims 31 and 38, the Examiner relies on modifying the method for configuration and management of storage resources of Nolan et al. with the method for handing

---

<sup>3</sup> See *Interconnect Planning Corp. v. Feil*, 227 USPQ 543 (CAFC 1985); see also *In re Fine*, 5 USPQ2d 1596, 1598 (CAFC 1988); see also *In re Gorman*, 18 USPQ 2d 1885, 1888 (CAFC 1991); see also *Al-Site Corp. v. VSI International, Inc.*, 50 USPQ2d 1161, 1171 (CAFC 1999).

<sup>4</sup> *In re Lee*, 61 USPQ2d 1430, 1433 (Fed. Cir. 2002); *In re Chu*, 36 USPQ2d at 1094.

off TCP sessions in a system of Nolan et al. to describe “the system wherein the computer is configured to maintain a pool of IP addresses and conform to the Dynamic Host Configuration Protocol (DHCP), and further wherein the NAS sensing module is configured to intercept a request from the remote storage device for an IP address” as recited by claims 31 and 38. Specifically, the Examiner refers to column 14 lines 60-65 and column 20 lines 58-67 of Lee et al. and stated that Lee et al. teaches “utilizing DHCP and IP address transporting.”

While Lee et al. does teach utilizing DHCP and IP address transporting, the Examiner’s assertion that it would have been *prima facie* obvious to one of ordinary skill in the art to use DHCP for the purpose of assigning and maintaining temporary host/device IP addresses is unreasonable and not supported by substantial evidence. The Examiner’s assertion is flawed as Nolan et al. fails to mention or even suggest detecting the attachment of a remote storage device as required by claims 31, and 38. Thus, modification of the method described by Nolan et al. with the method described by Lee et al. would fail to achieve the Applicant's claimed invention. Moreover, the Examiner has not pointed to any motivation in the art to modify the method described in Nolan et al. with the method described by Lee et al.

#### *Claims 34 and 35*

In rejecting claims 34 and 35, the Examiner proposes to modify the storage area network described by Nolan et al. with the master/slave architecture described by Lee et al. Specifically, the Examiner refers to column 17 lines 16-21 and column 20 lines 30-57 of Lee et al. and stated that Lee et al. teaches “a NAS master/slave architecture with a client coupled via a network.”

Neither Nolan et al. nor Lee et al. teach or suggest a NAS master configured to automatically incorporate storage capacity of the storage device into a logical volume when the NAS storage device is attached to the first network. As argued previously, Nolan et al. describes a storage domain manager that receives and manages storage transactions according to storage domains. Lee et al. generally relates to transparent access to network attached devices and does not describe automatically incorporating storage capacity whatsoever. Consequently, Lee et al. does nothing to overcome the deficiencies of Nolan et al. with respect to a NAS master remotely coupled to a NAS storage device and configured to automatically incorporate the storage capacity of the storage device.

Additionally, neither Nolan et al. nor Lee et al. teach or even suggest automatically incorporating storage capacity when the storage device is attached to the first network. For clarification purposes, Applicant has amended independent claim 34 to include subject matter recited in independent claims 1, 11, and 23. In other words, Applicant has amended independent claim 34 to require automatically incorporating storage capacity of the storage device upon detecting attachment of the storage device to the first network.

The Examiner's assertion that it would have been *prima facie* obvious to one of ordinary skill in the art to modify the storage area network described by Nolan et al. with the master/slave architecture described by Lee et al. is not supported by substantial evidence. The Examiner has not pointed to any motivation in the art for such a modification. The Examiner asserted that it would have been *prima facie* obvious to one of ordinary skill in the art to use a NAS master for the purpose of implementing a serverless network attached storage. This assertion is flawed as the motivation is not found within the art, as is required.

#### *Claim 36*

In rejecting claim 36, the Examiner proposes to modify the storage area network described in Nolan et al. with the master/slave architecture described by Lee et al. Specifically, the Examiner refers to column 2 lines 35-47 and column 23 line 54 – column 24 line 2 and stated that Lee et al. teaches an “inherent characteristic of using a NAS device, which is increasing storage space.”

Neither Nolan et al. nor Lee et al. teach or suggest a NAS master configured to automatically expand a size of a storage space mapped to the logical volume as recited by claim 36. As argued previously, Nolan et al. describes a storage domain manager that receives and manages storage transactions according to storage domains and does not mention automatically expanding a size of a storage space whatsoever. While Lee et al. describes expanding a size of a storage space, Lee et al. fails to mention automatically expanding a size of a storage, much less upon detecting attachment of a storage device to the network as required by independent claim 34 as amended, on which claim 36 is dependent. Consequently, Lee et al. does nothing to overcome the deficiencies of Nolan et al. with respect to a NAS master configured to automatically expand a size of a storage space mapped to the logical volume.

The Examiner's assertion that it would have been *prima facie* obvious to one of ordinary skill in the art to modify the storage area network described by Nolan et al. with the master/slave architecture described by Lee et al. is not supported by substantial evidence. The Examiner has not pointed to any motivation in the art for such a modification. The Examiner asserted that it would have been *prima facie* obvious to one of ordinary skill in the art to configure a NAS master automatic expansion for the purpose of accommodating and allocating the additional storage capabilities and resources. This assertion is flawed as the motivation is not found within the art, as is required.

*Claim 37*

In rejecting claim 37 the Examiner proposes to modify the storage area network described in Nolan et al. with the master/slave architecture described by Lee et al. Specifically, the Examiner refers to column 16 line 55 – column 17 line 22 and column 17 line 61 – column 18 line 21 and stated that Lee et al. teaches the use of “policy servers for storage incorporation.”

Neither Nolan et al. nor Lee et al. teach or suggest a policy manager to define one or more policies for incorporating the storage capacity as recited by claim 37. Nolan et al. describes configuration and management of storage transactions such as hot copy processes that migrate data between storage devices within a network. Lee et al. fails to describe a policy manager and does not mention policies for incorporating storage capacity whatsoever. Consequently, Lee et al. does nothing to overcome the deficiencies of Nolan et al. with respect to a policy manager that defines one or more policies for incorporating storage capacity as recited by claim 37.

The Examiner's assertion that it would have been *prima facie* obvious to one of ordinary skill in the art to modify the storage area network described by Nolan et al. with the master/slave architecture described by Lee et al. is not supported by substantial evidence. The Examiner has not pointed to any motivation in the art for such a modification. The Examiner asserted that it would have been *prima facie* obvious to one of ordinary skill in the art to configure a NAS master to comprise a policy manager for the purpose of maintenance of an efficient uniform file system. This assertion is flawed as the motivation is not found within the art, as is required.

For at least these reasons, the Examiner has failed to establish a *prima facie* case for non-patentability of Applicant's claims 8, 19, and 30-38 under 35 U.S.C. 103(a). Withdrawal of this rejection is requested.

### CONCLUSION

All claims in this application are in condition for allowance. Applicant respectfully requests reconsideration and prompt allowance of all pending claims. Please charge any additional fees or credit any overpayment to deposit account number 09-0069. The Examiner is invited to telephone the below-signed attorney to discuss this application.

Date:

10/21/4

Imation Legal Affairs  
P.O. Box 64898  
St. Paul, Minnesota 55164-0898  
Telephone: (651) 704-3604  
Facsimile: (651) 704-5951

By:

Eric D. Levinson

Name: Eric D. Levinson  
Reg. No.: 35,814

**AMENDMENTS TO THE DRAWINGS**

Applicant submits herewith a replacement drawing sheet for FIG. 1 for the present application. In particular, the reference character corresponding to the network by which storage users 8A-8N communicate with NAS storage devices 6A-6N has been changed from "8" to "10". Applicant has amended the second paragraph of page 4 lines 5-14, the third paragraph on page 4 lines 15-22, the fifth paragraph on page 4 lines 30-31 and first paragraph on page 5 lines 1-11, and the second paragraph on page 5 lines 12-19 of the present application as described in greater detail below to support the drawing change. No new subject matter has been added by way of this amendment.



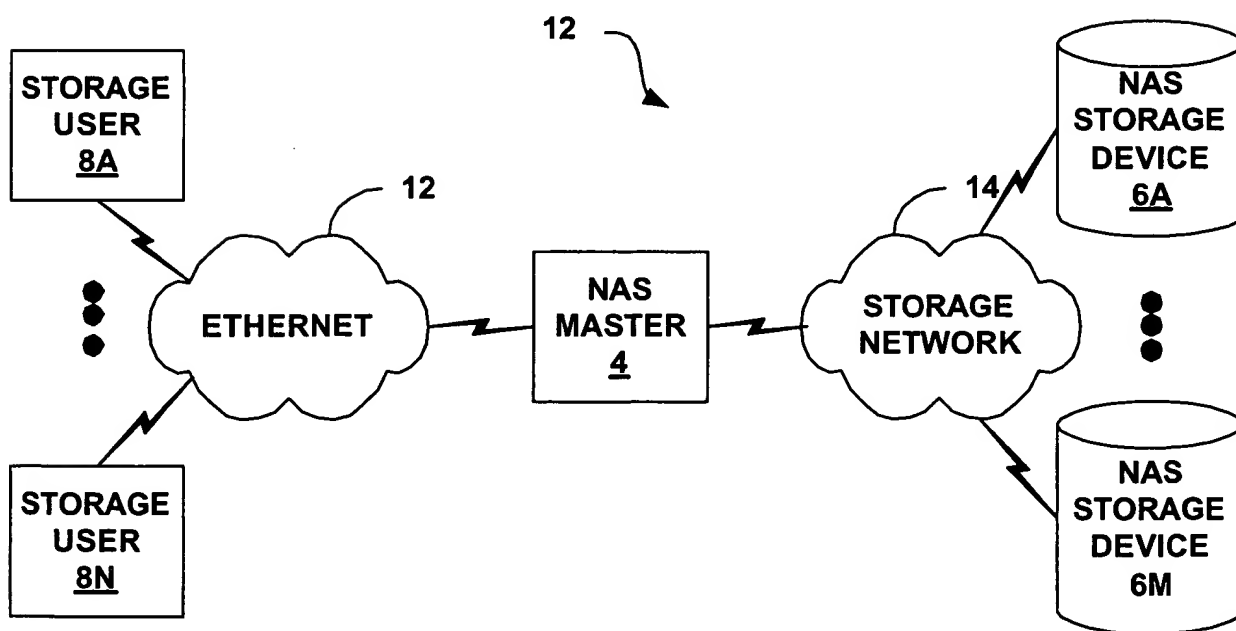



FIG. 2